

Optimal Control of Velocity and Nonlocal Interactions in the Mean-Field Kuramoto Model

Carlo Sinigaglia, Francesco Braghin, and Spring Berman

Abstract—In this paper, we investigate how the self-synchronization property of a swarm of Kuramoto oscillators can be controlled and exploited to achieve target densities and target phase coherence. In the limit of an infinite number of oscillators, the collective dynamics of the agents' density is described by a mean-field model in the form of a nonlocal PDE, where the nonlocality arises from the synchronization mechanism. In this mean-field setting, we introduce two space-time dependent control inputs to affect the density of the oscillators: an angular velocity field that corresponds to a state feedback law for individual agents, and a control parameter that modulates the strength of agent interactions over space and time, i.e., a multiplicative control with respect to the integral nonlocal term. We frame the density tracking problem as a PDE-constrained optimization problem. The controlled synchronization and phase-locking are measured with classical polar order metrics. After establishing the mass conservation property of the mean-field model and bounds on its nonlocal term, a system of first-order necessary conditions for optimality is recovered using a Lagrangian method. The optimality system, comprising a nonlocal PDE for the state dynamics equation, the respective nonlocal adjoint dynamics, and the Euler equation, is solved iteratively following a standard Optimize-then-Discretize approach and an efficient numerical solver based on spectral methods. We demonstrate our approach for each of the two control inputs in simulation.

I. INTRODUCTION

The emergence of collective behaviors from local interactions is ubiquitous in nature. Large groups of individuals can achieve complex macroscopic dynamical patterns by exploiting local interactions. The exact mechanisms by which flocks of birds and schools of fish respond to neighbors' behavior are still being investigated, although various models have been proposed in mathematical biology to achieve such macroscopic behaviors in simulation that resemble their natural counterparts in a qualitative way; see, e.g., the Vicsek model [1]. Inspired by such biological systems, the robotics community has devised control strategies for large-scale collectives of robots with the goal of mimicking the resilience, efficiency, and adaptive behavior of biological swarms [2]. Under the assumption of identical dynamics among individuals and a sufficiently large number of agents, it is possible to model the macroscopic behavior of a swarm of agents using a mean-field model [3], which describes the

This work was supported by the Italian Ministry of Education, University and Research (MIUR).

C. Sinigaglia and F. Braghin are with the Department of Mechanical Engineering, Politecnico di Milano, Milano, Italy, 20156 {carlo.sinigaglia, francesco.braghin}@polimi.it. S. Berman is with the School for Engineering of Matter, Transport and Energy, ASU, Tempe, AZ 85287 {spring.berman}@asu.edu.

density of agents at a specific location in the state space at a certain time. Mean-field models in the form of Louiville [4], Fokker-Planck [5], and McKean-Vlasov [6] equations are common in statistical physics and computational biology; they provide a powerful framework for studying the behavior of large-scale dynamical systems and have been recently analyzed from a control-theoretic perspective.

The steady-state collective behavior of such large-scale interacting systems is a prolific area of research in mathematical physics, along with the characterization of phase transitions in these systems with respect to a set of physical parameters. Very recently, the emergence of spatially inhomogeneous chimera states has been investigated for nonlocally interacting chiral particles [7], while the effect of non-reciprocal interactions between different species of individuals has been studied in [8].

In this paper, we investigate how interactions, synchronization patterns, and phase transitions can be exploited from a control perspective. We consider a large-scale swarm of identical Kuramoto oscillators, in which two possible control mechanisms affect the dynamics of each individual oscillator: an angular velocity control input and a control input that locally modulates the strength of oscillator interactions. We formulate an Optimal Control Problem (OCP) subject to the mean-field nonlocal Partial Differential Equation (PDE) model that describes the dynamics of the oscillator density, and we derive a set of first-order necessary conditions for optimality for this system. We then show in simulation how the control inputs can guide the density of oscillators to target distributions and accelerate their synchronization.

II. STATE DYNAMICS

The Kuramoto oscillator network [9] is a classical model that describes synchronization in a variety of biological and engineering systems. It consists of an array of coupled oscillators that each affect the others' phases via pairwise interactions. We consider a swarm of N oscillators in which the time-dependent phase $\theta_i(t)$ of each oscillator i is governed by the following controlled dynamical model:

$$\dot{\theta}_i = \omega_i + u_1 + \frac{u_2}{N} \sum_{j=1}^N \sin(\theta_j - \theta_i - \alpha) + \eta(t), \quad i = 1, \dots, N,$$

where ω_i is the oscillator's natural frequency, α is a constant phase shift that gives rise to traveling wave solutions [7], $\eta(t)$ is a white noise stochastic process whose intensity is associated with diffusion coefficient $D > 0$, u_1 is an angular velocity control input, and u_2 is a control gain that modulates

the strength of agent interactions. The usual Kuramoto phase model [9] is recovered by setting $u_1 = 0$ and $u_2 = K$, where K is a possibly non-constant interaction strength. Depending on the sign and magnitude of K and its dependence on θ , a variety of oscillator phase distributions, and hence degrees of synchronization, can appear [10], [11].

Kuramoto oscillators with noise are known to exhibit synchronization once the quantity $\frac{K}{D}$ exceeds a threshold, assuming that $K > 0$. Furthermore, traveling wave solutions and spatially inhomogeneous phases known as chimera states are known to arise for particular combinations of the interaction kernel and the phase shift α [12], [11], [10]. The complex polar order parameter defined below is used to measure phase coherence and synchronization intensity:

$$R(t)e^{i\psi(t)} = \frac{1}{N} \sum_{j=1}^N e^{i\theta_j(t)} \quad (1)$$

The amplitude R takes values in the interval $[0, 1]$, where 0 corresponds to disordered motion (incoherence) and 1 to global synchronization (phase locking). The variable ψ denotes the mean-field phase.

For the sake of simplicity, we assume that all oscillators have the same natural frequency, $\omega_i = \omega$. Without loss of generality, we can select $\omega = 0$, which is equivalent to expressing the oscillator dynamics in a frame rotating at angular velocity ω .

In the limit of an infinite number of oscillators, i.e., $N \rightarrow \infty$, the oscillator density $q(\theta, t)$ satisfies the following mean-field nonlocal PDE,

$$q_t - Dq_{\theta\theta} + \partial_\theta (u_2 w[q]q + u_1 q) = 0, \quad (2)$$

where

$$w[q] = \int_{S^{1'}} \sin(\theta' - \theta - \alpha) q(\theta', t) d\theta' \quad (3)$$

is a nonlocal integral term describing the oscillator interactions. Note that we use the notation $S^{1'}$ for the unit circle to highlight the role of the integration variable θ' .

Eq. (2) is a nonlocal parabolic PDE that describes the evolution in space-time of the probability density of a single oscillator having phase θ at time t , or equivalently, the normalized concentration of a population of oscillators at this phase and time. The natural functional space for q is the space of absolutely continuous probability measures, that is, $q \in V$ where $V = \{f \in L^2(0, T, H^1(S^1)) : \int_{S^1} f(\theta, t) d\theta = 1 \text{ and } f \geq 0 \text{ a.e. } t \in (0, T)\}$. In the following, we will refer to the solution of Eq. (2) with $u_1 \equiv 0$ and $u_2 \equiv K$ as the *uncontrolled mean-field dynamics*. In this case, we will set $K = 1$.

Equation (2) is defined on the unit circle S^1 and is thus naturally equipped with periodic boundary conditions. The mean-field equivalent [7] of the complex polar order parameter is:

$$R(t)e^{i\psi(t)} = \int_{S^1} e^{i\theta} q(\theta, t) d\theta \quad (4)$$

This continuous version of the polar order parameter has the same range for R as its discrete counterpart, since the oscillator density q is normalized to have unitary total mass.

In the following, we will make frequent use of Green's theorem for functions defined on S^1 . Given functions $f, g : S^1 \rightarrow \mathbb{R}$, we have that:

$$\int_{S^1} f_\theta g d\theta = - \int_{S^1} fg_\theta d\theta. \quad (5)$$

We will also need the following identity,

$$\int_{S^1} w[f]g d\theta = \int_{S^1} w^*[g]f d\theta, \quad (6)$$

where the functional w^* is defined as:

$$w^*[g] = \int_{S^{1'}} \sin(\theta - \theta' - \alpha) g(\theta') d\theta'.$$

It is straightforward to prove Eq. (6) by exchanging the variables of integration.

It is also easy to show that Eq. (2) conserves the total mass, $m(t) = \int_{S^1} q d\theta$, for any choice of control functions u_1, u_2 and parameters D, α . First, note that Eq. (2) can be equivalently written as follows, where $-Dq_{\theta\theta}$ is written as $\partial_\theta(-Dq_\theta)$:

$$q_t + \partial_\theta (u_2 w[q]q + u_1 q - Dq_\theta) = 0. \quad (7)$$

Then, we have that

$$\dot{m} = \int_{S^1} q_t d\theta = - \int_{S^1} \partial_\theta (u_2 w[q]q + u_1 q - Dq_\theta) d\theta = 0,$$

where we have substituted the state dynamics (7) and applied Eq. (5) with $g \equiv 1$. Furthermore, the velocity $w[q]$ of the transport field $w[q]q$, which is induced by the *nonlocal* wave, is bounded by 1, as we demonstrate in the next lemma.

Lemma II.1. *For all probability densities $q(t) \in H^1(S^1)$, $\|w[q(t)]\|_{L^\infty(S^1)} \leq 1$ for a.e. $t \in (0, T)$*

Proof. Using Eq. (3) and the fact that q is a probability density, we obtain the following inequality:

$$\begin{aligned} |w[q(t)]| &= \left| \int_{S^{1'}} \sin(\theta' - \theta - \alpha) q(\theta', t) d\theta' \right| \\ &\leq \sup_{\theta' \in S^{1'}} (|\sin(\theta' - \theta - \alpha)|) \int_{S^{1'}} |q| d\theta' \leq 1. \end{aligned}$$

Consequently, we have that

$$\|w[q(t)]\|_{L^\infty(S^1)} = \sup_{q(t) \in H^1(S^1)} |w[q(t)]| \leq 1.$$

□

In the following lemma, we use Lemma II.1 to obtain an upper bound on the norm of the mean-field state q . For this purpose, we select as control space $\mathcal{U} = L^2(0, T, L^\infty(S^1))$

Lemma II.2. *For any $u_1, u_2 \in \mathcal{U}$, it holds that*

$$\|q\|_V^2 \leq C_{\mathcal{U}}(u_1, u_2) \|q_0\|_{L^2(S^1)}^2, \quad (8)$$

where $q_0 \in L^2(S^1)$ is the initial probability density and $C_{\mathcal{U}}(u_1, u_2) \equiv \|u_1\|_{\mathcal{U}} + \|u_2\|_{\mathcal{U}}$.

Proof. We just need to show that Gårding's inequality holds for this problem, and then we can proceed as in Theorem 1 in [13]. The weak form of Eq. (2) can be written as:

$$\int_{S^1} q_t(t)\phi \, d\theta + a(q(t), \phi) = 0 \quad \forall \phi \in H^1(S^1)$$

for a.e. $t \in (0, T)$, where the form $a(q, \phi)$, defined as

$$a(q, \phi) = \int_{S^1} (Dq_\theta \phi_\theta - (u_2 w[q] + u_1)q \phi_\theta) \, d\theta,$$

is obtained by multiplying Eq. (2) by a test function $\phi \in H^1(S^1)$, integrating the equation over the domain S^1 , and then applying Eq. (5). We need to show that for almost every $t \in (0, T)$, there exists $\lambda(t) > 0$ such that:

$$a(q(t), q(t)) + \lambda(t) \|q(t)\|_{L^2(S^1)}^2 \geq \alpha_0(t) \|q(t)\|_{H^1(S^1)}^2$$

for some $\alpha_0(t) \geq 0$; i.e., that the form a is weakly coercive. First, by applying Lemma II.1, the Cauchy-Schwarz inequality, and Cauchy's inequality with ϵ (see e.g. [14], Appendix B.2), we can show that for every $\epsilon \geq 0$,

$$\begin{aligned} & \left| \int_{S^1} (u_2 w[q] + u_1)q \, q_\theta \, d\theta \right| \\ & \leq \| (u_2 w[q] + u_1) \|_{L^\infty(S^1)} \|q\|_{L^2(S^1)} \|q_\theta\|_{L^2(S^1)} \\ & \leq \left(\|u_2\|_{L^\infty(S^1)} + \|u_1\|_{L^\infty(S^1)} \right) \|q\|_{L^2(S^1)} \|q_\theta\|_{L^2(S^1)} \\ & = C \|q\|_{L^2(S^1)} \|q_\theta\|_{L^2(S^1)} \\ & \leq \frac{C^2}{4\epsilon} \|q\|_{L^2(S^1)}^2 + \epsilon \|q_\theta\|_{L^2(S^1)}^2, \end{aligned}$$

where $C \equiv C_{L^\infty(S^1)}(u_1, u_2) \equiv \|u_1\|_{L^\infty(S^1)} + \|u_2\|_{L^\infty(S^1)}$. Therefore, for a.e. $t \in (0, T)$, we can now write:

$$\begin{aligned} & a(q, q) + \lambda \|q\|_{L^2(S^1)}^2 \\ & \geq D \|q_\theta\|_{L^2(S^1)}^2 + \lambda \|q\|_{L^2(S^1)}^2 \\ & \quad - \left| \int_{S^1} (u_2 w[q] + u_1)q \, q_\theta \, d\theta \right| \\ & \geq (D - \epsilon) \|q_\theta\|_{L^2(S^1)}^2 + \left(\lambda - \frac{C^2}{4\epsilon} \right) \|q\|_{L^2(S^1)}^2 \\ & \geq \frac{D}{2} \|q_\theta\|_{L^2(S^1)}^2 + \left(\lambda - \frac{C^2}{2D} \right) \|q\|_{L^2(S^1)}^2 \\ & \geq \frac{D}{2} \|q_\theta\|_{L^2(S^1)}^2 + \frac{C^2}{2D} \|q\|_{L^2(S^1)}^2 \\ & \geq \alpha_0 \|q\|_{H^1(S^1)}^2, \end{aligned}$$

where we have chosen $\epsilon = \frac{D}{2}$, $\lambda = \frac{C^2}{D}$, and $\alpha_0 = \min\{\epsilon, \frac{\lambda}{2}\}$. Note that the time dependence of variables in the inequalities above is omitted to improve readability.

We can thus conclude that the form a is weakly coercive and apply the same reasoning as in [13], Theorem 1, to establish a bound on the norm of the mean-field state for every $u_1, u_2 \in \mathcal{U}$. \square

III. THE OPTIMAL CONTROL PROBLEM

In this section, we formulate an OCP subject to the mean-field dynamics in Eq. (2). A natural objective is to compute the minimum-energy control inputs u_1, u_2 that drive the

oscillator density to a target density $z(\theta, t)$ at a final time $t = T$. We define the cost functional to be minimized as

$$J(q, u_1, u_2) = J_q(q) + J_u(u_1, u_2),$$

where

$$\begin{aligned} J_q(q) &= \frac{\alpha_r}{2} \int_0^T \int_{S^1} (q(\theta, t) - z(\theta, t))^2 \, d\theta \, dt \\ &\quad + \frac{\alpha_t}{2} \int_{S^1} (q(\theta, T) - z(\theta, T))^2 \, d\theta, \end{aligned}$$

$$J_u(u_1, u_2) = \frac{1}{2} \int_0^T \int_{S^1} (\beta_1 u_1(\theta, t)^2 + \beta_2 u_2(\theta, t)^2) \, d\theta \, dt,$$

and $\alpha_r, \alpha_t, \beta_1$, and β_2 are nonnegative weighting constants.

The OCP can be written as:

$$\begin{aligned} \min_{u_1, u_2, q} \quad & J(q, u_1, u_2) = J_q(q) + J_u(u_1, u_2) \\ & q_t - Dq_{\theta\theta} + \partial_\theta (u_2 w[q]q + u_1 q) = 0, \\ \text{s.t.} \quad & q(\theta, 0) = q_0(\theta). \end{aligned} \quad (9)$$

We consider the control inputs u_1 and u_2 to be acting simultaneously in the OCP, and we derive a set of first-order necessary optimality conditions accordingly using a Lagrangian method. In order to do so, we define the Lagrangian functional [15] as:

$$\begin{aligned} \mathcal{L}(q, u, p) &= J(q, u_1, u_2) \\ &\quad - \int_0^T \int_{S^1} (q_t - Dp_{\theta\theta} + \partial_\theta (u_2 w[q]q + u_1 q)) p \, d\theta \, dt, \end{aligned}$$

where $p : S^1 \times (0, T) \mapsto \mathbb{R}$ is the adjoint field still to be determined. The adjoint dynamics are recovered by taking the Gateaux derivative with respect to the state q and setting it to zero for any state variation. We focus on deriving the first variation of the third term of the PDE constraint in (9), since the treatment of the time derivative and diffusion terms is standard in the literature on OCPs for parabolic equations (see e.g. [15]). First, it is useful to note that:

$$\begin{aligned} & \int_0^T \int_{S^1} \partial_\theta (u_2 wq) p \, d\theta \, dt \\ &= - \int_0^T \int_{S^1} u_2 w[q] q p_\theta \, d\theta \, dt. \end{aligned} \quad (10)$$

Taking the first variation of the term of interest, we have:

$$\begin{aligned} & \frac{\partial}{\partial \epsilon} \Big|_0 \int_{S^1} -u_2 wq + \epsilon \psi p_\theta \, d\theta \, dt \\ &= \int_{S^1} -u_2 (w[q]\psi + w[\psi]q) p_\theta \, d\theta \, dt \\ &= - \int_{S^1} (u_2 w[q]p_\theta + w^*[u_2 p_\theta q]) \psi \, d\theta \, dt, \end{aligned}$$

where we have used Eq. (6) to take out ψ from the nonlocal operator w . Regarding the component of the term that contains u_1 , we use Eq. (5) to write:

$$\int_{S^1} \partial_\theta (u_1 q) p \, d\theta = - \int_{S^1} u_1 q p_\theta \, d\theta$$

so that the state variation of this term reads:

$$\frac{\partial}{\partial \epsilon} \Big|_0 \int_{S^1} u_1 (q + \epsilon \psi) p_\theta d\theta = \int_{S^1} u_1 p_\theta \psi d\theta. \quad (11)$$

The adjoint dynamics can be then written as:

$$-p_t - D p_{\theta\theta} - (u_2 w[q] + u_1) p_\theta - w^* [u_2 p_\theta q] = \alpha_r (q - z), \quad (12)$$

with the final time condition specified as:

$$p(\theta, T) = \alpha_t (q(\theta, T) - z(\theta, T)).$$

We can now derive the Euler equation for the reduced gradient with respect to the controls u_1 and u_2 . Using Eq. (11), the first variation of the Lagrangian with respect to the velocity control input u_1 is straightforward to compute:

$$\nabla J_{u_1}(\theta, t) = \beta_1 u_1(\theta, t) + q(\theta, t) p_\theta(\theta, t). \quad (13)$$

This control mechanism for mean-field PDEs, i.e., optimal control of a velocity field, has been considered in various applications (see e.g. [13], [5]), and controllability results [16] have been obtained for the case where agent interactions are not present (i.e., $u_2 \equiv 0$). In order to recover the reduced gradient with respect to u_2 , we make use of Eq. (10), which is linear in u_2 . We can then find that the reduced gradient has the form:

$$\nabla J_{u_2}(\theta, t) = \beta_2 u_2(\theta, t) + w[q](\theta, t) q(\theta, t) p_\theta(\theta, t). \quad (14)$$

Note that the integral term $w[q]$ makes this gradient equation nonlocal, in contrast to gradient equation (13), which means that at each point in space-time, the equation depends on the entire state solution q .

IV. SIMULATION RESULTS AND DISCUSSION

In this section, we solve the OCP numerically using an iterative gradient descent method and analyze the results. At each iteration, the current control input values u_1 , u_2 are used to solve the state and adjoint equations, Eqs. (2) and (12). The reduced gradient is then computed using Eqs. (13) and (14), and the control inputs are updated such that they produce a sufficient decrease in the cost functional. The numerical algorithm is implemented in Python and the operators are discretized using the open-source software Dedalus [17], which makes efficient use of novel spectral algorithms for the solution of PDEs. Although we have derived optimality conditions for the general case where both u_1 and u_2 are applied, we optimize u_1 and u_2 separately in the simulations in order to investigate the ability of each input to steer the state dynamics towards the target density.

In the u_1 -controlled case, we set $u_2 \equiv K = 1$ and design the velocity control input $u_1(\theta, t)$ to steer an initial density of oscillators towards a target density $z(\theta)$ at final time T . A similar control problem is solved in [5], which does not include agent interactions, and in [6], which includes interactions in the form of spatially attractive/repulsive potentials. Our aim is to investigate the effect of the control input u_1 on the synchronization mechanism, as compared to the uncontrolled dynamics ($u_1 = 0, u_2 = 1$).

In the u_2 -controlled case, we set $u_1 = 0$ and design the interaction strength control input $u_2(\theta, t)$, which locally regulates the intensity of agent interactions, to solve this density steering problem. At each iteration of the optimization algorithm, $u_2(\theta, t)$ is updated using Eq. (14); the adjoint equation (12) is the same for both the u_1 -controlled and u_2 -controlled cases. We note that the influence of the integral term $w[q]$ may limit the control authority with respect to u_2 ; these restrictions on controllability must be further investigated.

In all simulations, we set $T = 10$ s and choose $D = 0.25$ and $\alpha = 0$, for which a synchronized phase is the solution of the uncontrolled system at steady-state [7] in the form of a Gaussian density that can be characterized semi-analytically. The objective of the OCP is to speed up the convergence of the state dynamics to the synchronized phase while steering the mean phase to the mean of the target density. Toward this end, we define the target density as a relatively low-variance Gaussian function with a mean of $\frac{3}{2}\pi$ rad.

Fig. 1 plots the time evolution of the phase coherence parameter, the amplitude of the complex polar order parameter in Eq. (4), for both controlled systems and the uncontrolled system. The trajectories of the phase coherence parameter and mean phase in the complex plane for all three systems are plotted in Fig. 2, along with the target coherence and mean phase. These figures show that both control inputs u_1 and u_2 are able to steer the system towards the target coherence and mean phase, with u_1 producing synchronization more quickly than u_2 . In contrast, the uncontrolled system achieves synchronization much more slowly, and it converges to a mean phase that differs from the target phase. Figure 5 plots the u_1 -controlled and u_2 -controlled densities and the uncontrolled density at the final time $T = 10$ s, along with the target density. The controlled densities have both closely approached the target density, while the uncontrolled system has not yet synchronized.

The state dynamics of the uncontrolled, u_1 -controlled, and u_2 -controlled systems over the entire simulation are shown in Figs. 3, 4, and 7, respectively. The u_2 -controlled system exhibits more complex dynamics than the u_1 -controlled system before converging to the target density; this is due to the previously mentioned limitations on control authority with respect to u_2 . The space-time evolution of the corresponding control input u_1 and the controlled nonlocal transport field $w[q]u_2$ are plotted in Figs. 6 and 8, respectively. The quantity $w[q]u_2$ is plotted rather than u_2 in order to directly compare the controlled transport fields produced by u_1 and u_2 .

It is worthwhile to comment on the signs of the control inputs. The positive (negative) sign of the velocity control input u_1 generates a transport field in the direction of increasing (decreasing) θ , while the role of the sign for the interaction strength control input u_2 is more subtle. Positive values of u_2 drive the corresponding oscillator density toward alignment, and thus synchronization, while negative values produce oscillator disalignment [11]. Thus, in order to reach a target density, both locally repelling and locally aligning interaction strength inputs are needed.

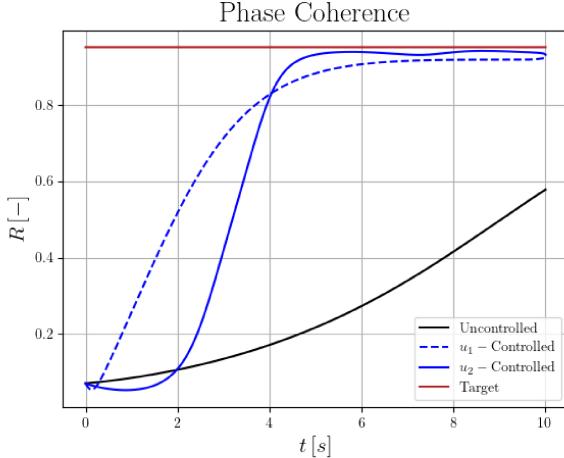


Fig. 1. Phase coherence parameter $R(t)$ over time. Both the velocity and interaction strength control inputs speed up the system self-synchronization.

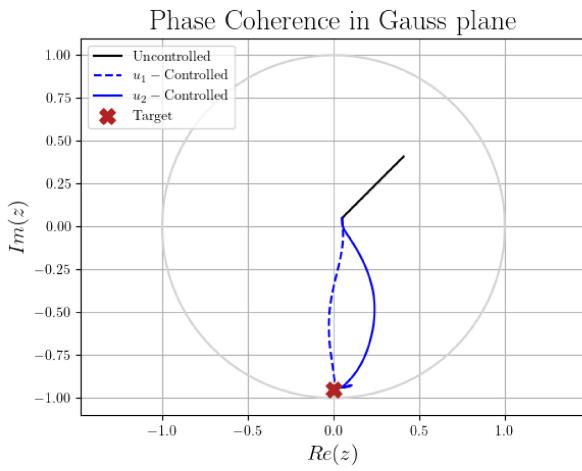


Fig. 2. Image of the continuous synchronization metric $R(t)e^{i\psi(t)} : [0, T] \mapsto \mathbb{C}$. Trajectories of both controlled systems converge to the target coherence and mean phase.

In summary, both the velocity and interaction strength control inputs are able to successfully drive the oscillator density to a target distribution while increasing the self-synchronization speed. Interaction strength control is able to achieve similar performance as velocity control, with slightly slower convergence to synchronization, despite its controllability limitations.

V. CONCLUSION

In this paper, we have presented a novel Optimal Control Problem (OCP) for a large-scale system of identical Kuramoto oscillators in the mean-field limit. The mean-field state dynamics in the OCP are nonlocal due to an integral term describing oscillator interactions; this term is shown to be a bounded functional, and some preliminary results on the well-posedness of the state dynamics are established. Both control mechanisms considered, the angular velocity and interaction strength of the oscillators, are able to accurately

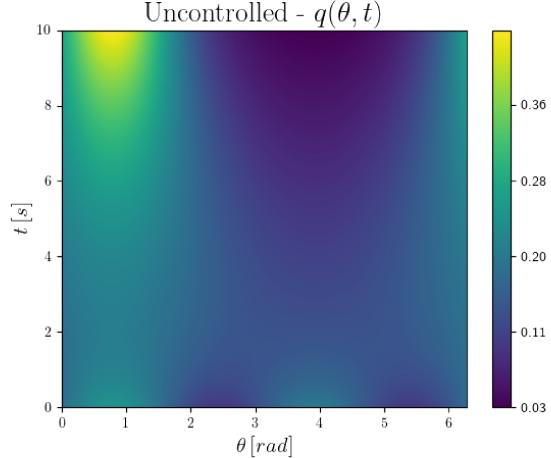


Fig. 3. Uncontrolled state dynamics with diffusion parameter $D = 0.25$ and interaction strength $K = 1$. Synchronization emerges close to the final time, with mean phase dependent on the initial conditions.

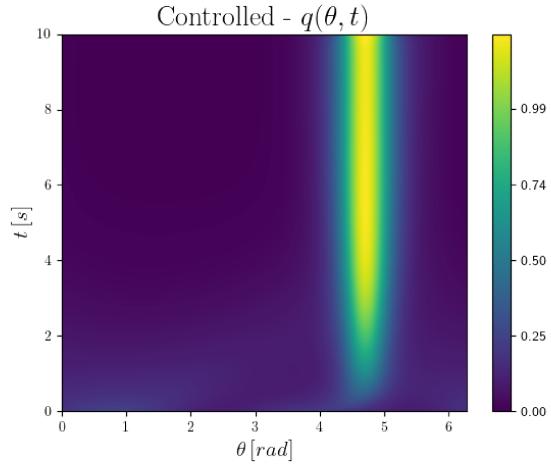


Fig. 4. Controlled state dynamics under optimal velocity control input u_1 . The density is coherent and synchronized after a brief transient period.

track a target density, as well as increase the synchronization speed over that of the uncontrolled dynamics. Furthermore, the interaction strength control input achieves the target density by exploiting the nonlocal term. These results set the grounds for establishing controllability and existence results for OCPs in which the state dynamics are governed by self-synchronizing nonlocal PDEs, which in turn are envisioned to find application in the control of collective systems such as robotic swarms through these synchronization mechanisms.

REFERENCES

- [1] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, “Novel type of phase transition in a system of self-driven particles,” *Physical review letters*, vol. 75, no. 6, p. 1226, 1995.
- [2] R. Pfeifer, M. Lungarella, and F. Iida, “Self-organization, embodiment, and biologically inspired robotics,” *Science*, vol. 318, no. 5853, pp. 1088–1093, 2007.
- [3] K. Elamvazhuthi and S. Berman, “Mean-field models in swarm robotics: A survey,” *Bioinspiration and Biomimetics*, vol. 15, no. 1, 2020.

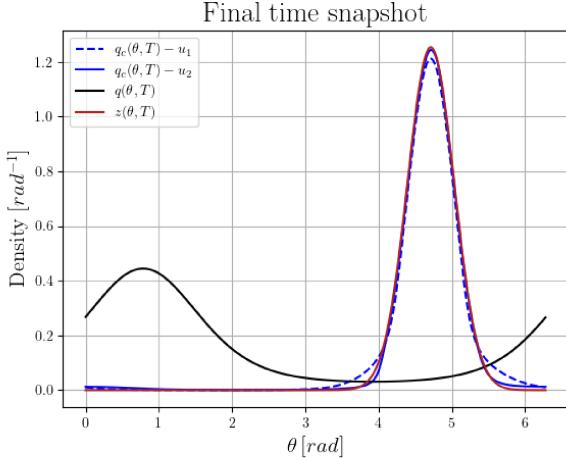


Fig. 5. Snapshot of the state $q(\theta, T)$ at final time $T = 10$ s. The uncontrolled system (black) has not reached a coherent phase, while the controlled systems (blue) are both close to the target density (red).

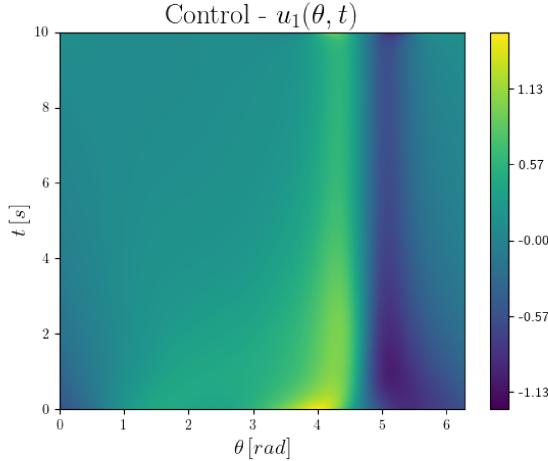


Fig. 6. Space-time evolution of the optimal velocity control input $u_1(\theta, t)$. Once the state $q(\theta, t)$ is steered close to the target density, the control magnitude is highest in the vicinity of the mean phase of the target density, inducing both positive and negative local transport fields.

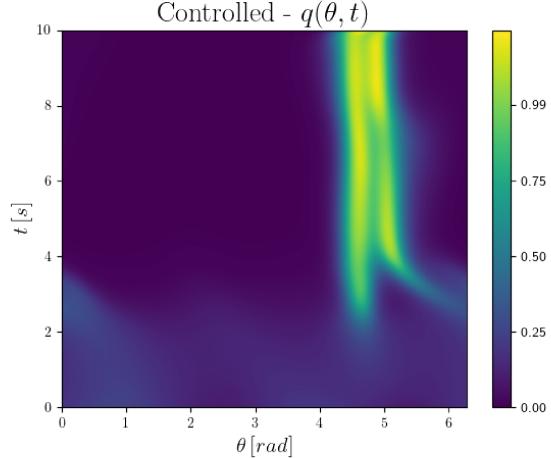


Fig. 7. Controlled state dynamics under optimal interaction strength control input u_2 . The transient period before synchronization accelerates is longer than in the u_1 -controlled case, and the steady-state density is less coherent.

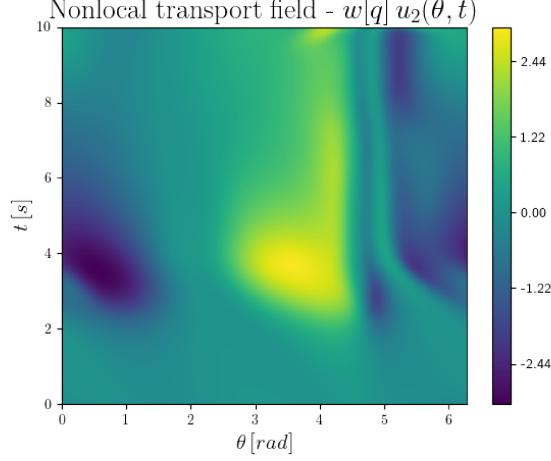


Fig. 8. Space-time evolution of the controlled nonlocal transport field $w[q]u_2$ induced by the optimal interaction strength control input $u_2(\theta, t)$. A negative control action is associated with oscillators that tend to disalign. The control intensity, given by u_2 (not shown), is approximately seven times higher than u_1 in the u_1 -controlled case due to the influence of the nonlocal term $w[q]$, which directly multiplies u_2 .

- [4] N. Pogodaev, “Optimal control of continuity equations,” *Nonlinear Differential Equations and Applications NoDEA*, vol. 23, 04 2016.
- [5] K. Elamvazhuthi, H. Kuiper, and S. Berman, “PDE-based optimization for stochastic mapping and coverage strategies using robotic ensembles,” *Automatica*, vol. 95, pp. 356 – 367, 2018.
- [6] Y. Chen, “Density control of interacting agent systems,” pp. 1–14, 2021. [Online]. Available: <http://arxiv.org/abs/2108.07342>
- [7] N. Kruk, J. A. Carrillo, and H. Koepl, “A finite volume method for continuum limit equations of nonlocally interacting active chiral particles,” *Journal of Computational Physics*, vol. 440, p. 110275, 2021.
- [8] M. Fruchart, R. Hanai, P. B. Littlewood, and V. Vitelli, “Non-reciprocal phase transitions,” *Nature*, vol. 592, no. 7854, pp. 363–369, 2021.
- [9] Y. Kuramoto, *Chemical Oscillations, Waves, and Turbulence*, ser. Dover Books on Chemistry Series. Dover Publications, 2003.
- [10] D. M. Abrams and S. H. Strogatz, “Chimera states for coupled oscillators,” *Physical Review Letters*, vol. 93, no. 17, pp. 1–4, 2004.
- [11] H. Hong and S. H. Strogatz, “Kuramoto model of coupled oscillators with positive and negative coupling parameters: An example of conformist and contrarian oscillators,” *Physical Review Letters*, vol. 106, no. 5, pp. 1–4, 2011.
- [12] H. Hong, K. P. O’Keeffe, and S. H. Strogatz, “Phase coherence induced by correlated disorder,” *Physical Review E*, vol. 93, no. 2, pp. 1–6, 2016.
- [13] C. Sinigaglia, A. Manzoni, and F. Braghin, “Density control of large-scale particles swarm through PDE-constrained optimization,” pp. 1–18, 2021. [Online]. Available: <http://arxiv.org/abs/2104.06373>
- [14] L. C. Evans, *Partial Differential Equations*. American Mathematical Society, 2010.
- [15] F. Tröltzsch, *Optimal Control of Partial Differential Equations*. American Mathematical Society, 2010, vol. 112.
- [16] K. Elamvazhuthi, H. Kuiper, M. Kawski, and S. Berman, “Bilinear controllability of a class of advection–diffusion–reaction systems,” *IEEE Transactions on Automatic Control*, vol. 64, no. 6, pp. 2282–2297, 2019.
- [17] K. J. Burns, G. M. Vasil, J. S. Oishi, D. Lecoanet, and B. P. Brown, “Dedalus: A flexible framework for numerical simulations with spectral methods,” *Physical Review Research*, vol. 2, no. 2, 2020.